Putting the odds on your side

In the first of three articles, Cynthia Kase shares some of her proprietary trading tools and explains how statistically based indicators can outperform their standard counterparts.

By Cynthia Kase

Many traditional technicians use indicators designed before the advent of the personal computer and are unaware that a new category of Wall Street professional, the financial engineer, is bringing the same analytical approaches to the market that physical scientists and engineers bring to the physical world.

Harnessing the power of high-speed PCs and using basic financial engineering principals, we've developed some new indicators. Here we review a statistically based momentum indicator and a highly accurate stop system. Next month we'll cover a statistical screening system and a replacement for the MACD that employs random walk theory. Finally, we'll show how to combine these techniques in a trading situation.

For whom the bell tolls Most traders are familiar with the normal bell curve distribution, but are unaware of the mathematics and probability theory underlying it. Data (including most equity and futures data) that is affected by a high number of minor and independent random events is, over time, roughly normally distributed. Thus, understanding the normal bell curve and stochastic (random) behavior is important for the technician.

The function for the bell curve is:

$$f(x\,|m,s) = (1/(s*\sqrt{(2\Pi)}))*e^{-0.5((x-s)/m)^2}$$

where m is the mean and s is the standard deviation (the square root of the variance).

In a simple distribution in which there are two possible outcomes, like a coin toss, the variance is:

$$\text{variance} = s^2 = n*p(1-p)$$

where n is the number of data points and p is the probability of one outcome.

In trading, outcomes are more complex. Unlike a coin toss, there are more than two possible prices. Here, the variance is the sum of the squared deviations between the data and the mean (as shown above) where x is the data point, and m is the mean.

$$\text{variance} = s^2 = 1/(n-1)\Sigma(x-m)^2$$

Let's assume we have a population of six basketball players and six small children with heights of: 7', 4'2", 7'1", 4'1", 7'2", 4'3", 6'11", 3'11", 7'2", 4'5", 6'11", and 3'11", for an average height of 5'7" or 67 inches.

Using the previous formula, the standard deviation of this population is the square root of $(1/11)*((84-67)^2 + (50-67)^2 + (x-m)^2$ etc.)), or 18.4". If we had a population of chorus girls, the shortest of whom is 5'5" and the tallest is 5'8", the variance and standard deviation would be much smaller, although the average height would be the same as the first group.

Now we can begin to understand some very important precepts about the market:

1. Random markets are mean-reverting (returning to some average value) and roughly form a normal distribution — the bell curve.
2. In such markets, the average daily rate of change is zero; any change in price results from the variation in the rate.
3. Volatility is related to the square root of time because standard deviation is the square root of variance.
4. Normally distributed data has definable probabilities. (If you know where a data point lies on the bell curve, you know the probability of its occurrence.)
5. A price out of the range of random market probabilities is said to result from "bias," i.e., trend.

"Random vs. biased" (below) shows a Latin Hypercube simulation illustrating these points. This type of simulation uses a random number generator and, given certain probability criteria, generates empirical results. Assume a fictitious contract with a price of \$5.50 per oz. and volatility of 30%. Using the Latin Hypercube simulation we can compare a random market with no bias to a biased market with a 0.4% increase per day to find the price after one quarter.

First let's employ rule three. The standard deviation of the daily rate of change is $30\%/\sqrt{255}$ (the approxi-

Random vs. biased

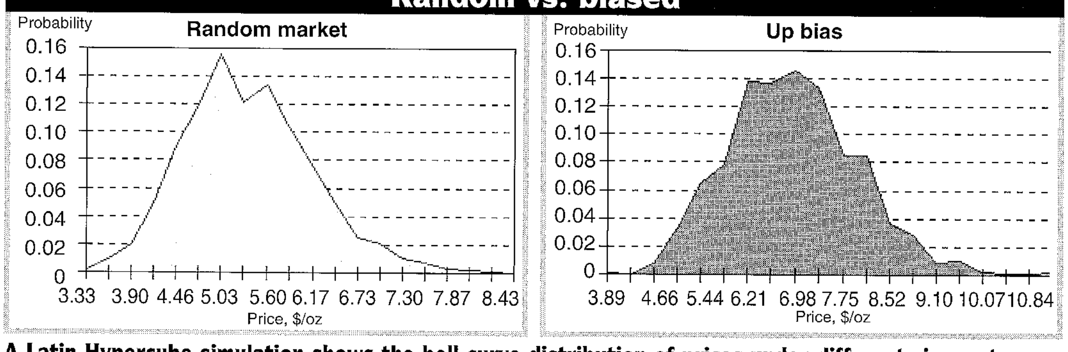

A Latin Hypercube simulation shows the bell curve distribution of prices under different circumstances: At left, a no bias situation; at right, a built-in increase of 0.4% per day.

Statistical silver

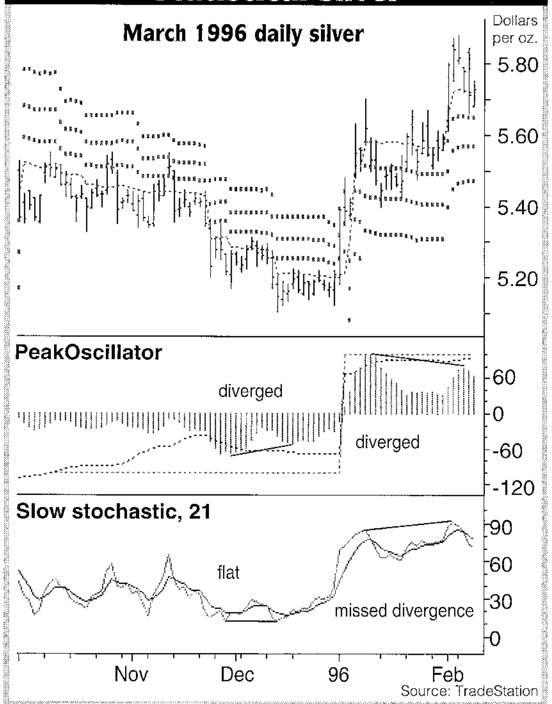

The PeakOscillator signals divergences other oscillators miss. Top: Dev-stop levels one, two and three.

mate number of trading days in a year), or 1.9%. The "Random market" and "Up bias" sections of the chart show the results. The first illustrates rules one, two and three. The resultant mean expectation is the starting price; one standard deviation is 78¢, or about 15% of the underlying contract price, which equals the annualized volatility divided by the square root of four — the expected theoretical quarterly volatility. In other words, the simulation and the theory produced the same result.

In the second chart, the most probable price is \$7.13. Elaborating on rule four, one standard deviation around the mean captures about 67% of all occurrences and two standard deviations about 95%. Thus, there is a direct relationship between where a data point falls on the curve and its probability. Statisticians consider any data outside two standard deviations improbable. Two standard deviations over the mean is roughly \$5.50 + (2*0.78), or \$7.06. The mean expectation price of \$7.13 is higher than the upper limit or "confidence level," so this market was, over the observed time horizon, biased or trending. These basic principles underlie the Random Walk Index (RWI), popularized by Mike Poulous and, in turn, a new momentum indicator — the PeakOscillator.

The RWI measures the degree to which price movement up or down differs from the random model:

$$RW_{high} = \frac{high_0 - low_n}{ATR * \sqrt{n}} \quad RW_{low} = \frac{high_n - low_0}{ATR * \sqrt{n}}$$

where *ATR* = Average True Range, n = number of bars in the "look back" period, and the RW itself is the maximum value achieved when n is run from 8 to 65.

Because a normal distribution requires a high number of data points, the distribution of average range is highly unstable for low "n" values. Therefore, we always use a minimum of 30 data points when calculating the ATR function.

Peak performance Simple momentum oscillators take the difference between two moving averages. To construct our new oscillator, we substitute the Random Walk Index with modified ATR and smooth the result. While simple in concept, a fair amount of computing capability is required to generate the PeakOscillator.

This indicator has a number of advantages over simple momentum indicators. The traditional stochastic and RSI are normalized to fall between 0 and 100, determined solely by the immediate price environment; they cannot be compared across different times or commodities (the RSI on a March 1996 daily silver chart cannot be compared directly to the RSI on a December 1986 hourly crude oil chart).

The volatility-adjusted PeakOscillator has no such limits. It is normalized for range, which is proportional in the shorter term to volatility; accordingly, the indicator is universal because it is adjusted for the "level playing field" of the market — volatility.

Given this universality, we can randomly sample data from many markets under many conditions, develop a distribution of the PeakOscillator, and directly determine the relationship between the indicator's value and the probability that higher or lower values will occur. We've evaluated the oscillator over more than 80 years of commodity prices and determined the 90th percentile.

Accordingly, the PeakOscillator is accompanied by two "extreme" lines: the local line, which is the 98th percentile of the local distribution of the oscillator, and the PeakOut line, which is the higher of the local line and the 90th percentile of the 80-year history. In cases where traditional divergence occurs, the next-to-last peak often will penetrate an extreme line, providing advance notice that a pullback may occur, followed by a new divergent price surge.

Another major advantage of the PeakOscillator is reliability. It generates divergence signals where traditional momentum indicators do not. "Statistical silver" shows two divergence signals in the March 1996 silver contract missed by the popular slow stochastic, with advance notice of the divergence from the penetration of the extreme lines — the local line at the lows and the PeakOut ➡

Catching spikes

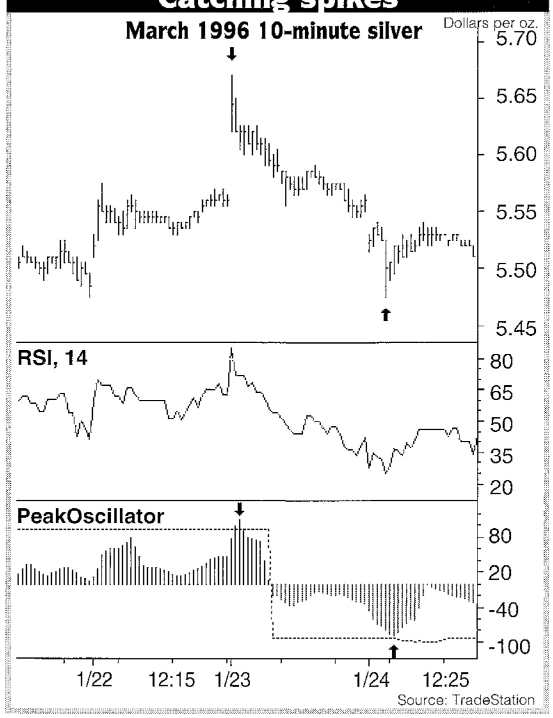

The PeakOscillator catches non-divergent turns: Note the penetration of the extreme line at the spikes.

line toward the highs.

The PeakOscillator also can signal a reversal when divergence isn't possible. Divergence requires a comparison between either two highs or lows. With a spike top or bottom, there is only one extreme. The PeakOscillator's penetration of the PeakOut line can signal an imminent turn where all other normalized indicators fail. "Catching spikes" (page 35) shows two spike-type turns signaled by the PeakOscillator but missed by the RSI.

The Dev-Stop A second indicator is the Dev-Stop, which uses standard deviations of range to set statistically based stop points. Many traders set stops based on what they can afford. Using Charles Dow's analogy of the market to an ocean, we know if we want to go sailing, we cannot demand the sea be calm and the swells small. We must be prepared to take the risks determined by the ocean, not vice versa. Similarly, we must accept the risk imposed by the market, not that dictated by our risk appetite — adjusting our time frame and volume if the risk in a given scenario is too high.

After reading Wilder's 1978 *New Concepts in Technical Trading Systems*, we modified his volatility system and used a fixed factor multiplied by the average true range (which is proportional to volatility) for a stop. However, we found this approach incomplete, as it neither considered the variance of volatility nor the tendency of volatility to spike, forming a right-skewed distribution rather than a normal one.

Using the earlier example of the two groups with the same average height of 5'7", a doorway that allows 98% of the first group to pass must be higher than one for the second group. Similarly, a market with a greater volatility must have wider stops than one which trades smoothly. Simply, the Dev-Stop is based on reversals of 1, 2.2 and 3.6 standard deviations over the mean of a two-bar true range against the trend. The extra 10% and 20% on the second and third standard deviations is to allow for the right-hand skew.

To calculate the stop, first take the average and standard deviations of a two-bar true range; next, add the deviation (multiplied by the appropriate factor) to the mean. The resultant reversal value is then subtracted from the highest profit point, if long, and added, if short. Three levels are used as scale-out bands, in addition to the mean, which is used as a warning.

The stop, as shown in "Statistical silver," is displayed by an underlying simple double moving average to default long or short. Notice how well Dev-Stop 3 held the trend both up and down, with the intermediate stops breaking on minor corrections. Generally, you would use level three, liquidating half your position and pulling the stop into level one on confirmed divergences.

Because there are only so many ways to finesse the front end of a trade, improving exit techniques is helpful in improving trading results. New PCs now allow for a more rigorous statistical approach and more accurate indicators. **FM**

Cynthia Kase is president of Kase and Co. in Albuquerque, N.M. and author of Trading With The Odds.